Learning Theory

CISC 5800
Professor Daniel Leeds

The classifier

Function C that provides
correct label (Y) based on features (X)

C(x)=y

lion: 16
wolf: 12
monkey: 14
broker: 0
analyst: 1
dividend: 1

$\xrightarrow{C}$ jungle

lion: 0
wolf: 2
monkey: 1
broker: 14
analyst: 10
dividend: 12

$\xrightarrow{C}$ wallStreet

Goal: identifier classifier that maximizes
correct labels for most inputs

Sample complexity

How many training examples needed to learn concept?

- *X* – set of data points
- *P(X)* – Probability of drawing data point x
- *H* – space of hypotheses H = {h : X -> classes }
- *C* – *correct* assignment C = {c : c(x) = y ∀xϵX }

Probability of error

X – data points

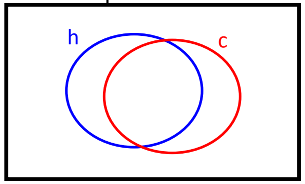

H = {h : X -> {0,1}}

True error of h: probability randomly selected data point from P(X) misclassified

$error_{true}(h) = Pr_{x \sim P(X)} [h(x) \neq c(x)]$

- Hard to compute, but can prove properties of $error_{true}$

Example: Learner picks one of fixed number of classifiers $h \epsilon H$

Correct classifier c is some assignment of each x to a label

How many training points m needed for $error_{true}(h) < \varepsilon$?

$Prob[error_{true}(h) \leq \varepsilon] > 1\text{-}\delta$

"Probability learned classifier h has worse than ε error is < δ"

"Probably Approximately Correct Learning" – PAC Learning

5

Binary example: sample complexity

Note for $\varepsilon = [0,1]$, $(1-\varepsilon) \leq e^{-\varepsilon}$

What is the chance learned h is bad but classifies training data correctly?

If $error_{true}(h) > \varepsilon$:

- Prob [h correctly labels x^1] < $(1-\varepsilon) \leq e^{-\varepsilon}$
- Prob [h correctly labels x^1 and x^2 … and x^m] < $(1-\varepsilon)^m \leq e^{-m\varepsilon}$

If classifier picks one h* randomly from H

- Prob[h^* is bad] = Prob[h_1 bad] + … Prob[h_n bad]
  = Prob[$error_{true}(h^*) > \varepsilon$] < $|\mathrm{H}|\ e^{-m\varepsilon}$

Valiant, 1984

6

Binary example: sample complexity

Number of data points to reduce chance of false classification, enforce

$Prob[error_{true}(h) \leq \varepsilon] > 1\text{-}\delta$

$1\text{-}\ Prob[error_{true}(h) \leq \varepsilon] = Prob[error_{true}(h) > \varepsilon] < \delta$

Prob[$error_{true}(h^*) > \varepsilon$] < $|\mathrm{H}|\ e^{-m\varepsilon}$; stricter bound $|\mathrm{H}|\ e^{-m\varepsilon} < \delta$

Valiant, 1984

8

Binary example: sample complexity

Number of data points to reduce chance of false classification, enforce

$Prob[error_{true}(h) \leq \varepsilon] > 1\text{-}\delta$

Prob[$error_{true}(h^*) > \varepsilon$] < $|\mathrm{H}|\ e^{-m\varepsilon} < \delta$

$$m > \frac{1}{\varepsilon} \ln \frac{|H|}{\delta}$$

Valiant, 1984

9

VC Dimensions

If H not finite, PAC result seems to require ∞ data points

- Overly conservative

“Dichotomy” – division of set of points S into two subsets

- “Shattering” – set of points is **shattered** by H iff there exists h∈H associated with every possible dichotomy

Vapnik-Chervonenkis dimension **VC(H)** is size of largest finite subset of S that can be shattered by H

10

Shattering example3

- H={rectangles: inside is 1, outside is 0} VC(3)
- S={3 specified dots}

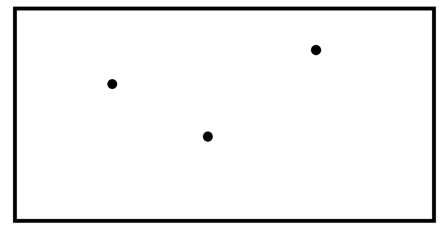

12

Shattering example

- H={rectangles, inside is 1 outside is 0} VC(3)
- S={4 specified dots}

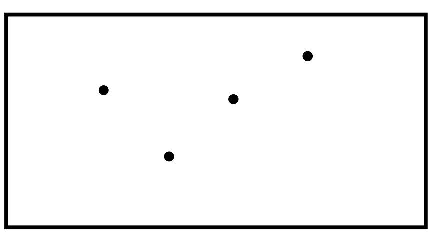

14

Shattering example

- H2={rectangles, inside is 1 outside is 0
  inside is 0 outside is 1}
  VC(4)
- S={4 specified dots}

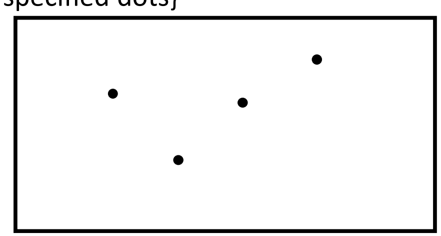

15

Shattering example

- H={rectangles, inside is 1 outside is 0} VC(4)
- S={4 specified dots}

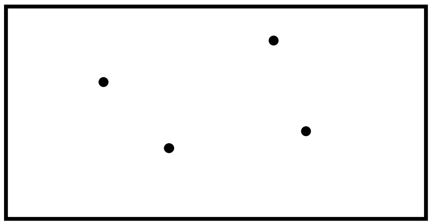

Shattering example 4

- H={rectangle, inside is 1 outside is 0}
- S={8 specified dots}

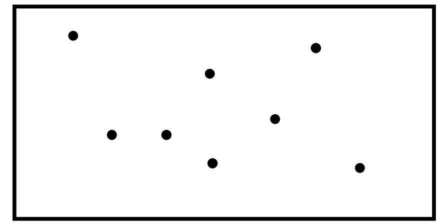

Shattering example 4

- H={rectangle, inside is 1 outside is 0} **H(4)**
- S={8 specified dots}

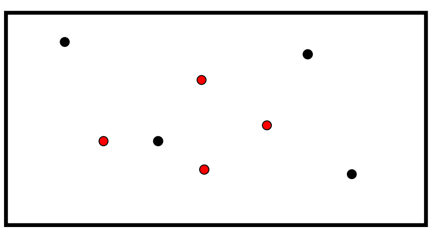

Shattering infinite points

- H={Linear separators}
- S={Any point in 2D feature space}
- S={Any point in nD space}

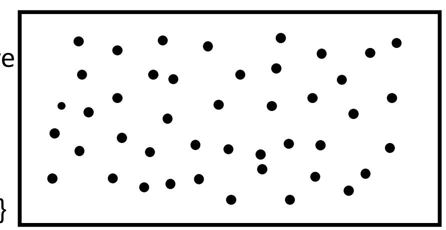

PAC result with infinite H

VC(H) is size of largest finite subset of X that can be shattered by H

- d=$VC(H)$
- $m \geq O\left(\frac{1}{\varepsilon}\left[d \log \frac{1}{\varepsilon} + \log \frac{1}{\delta}\right]\right) \sim \frac{1}{\varepsilon}\left[d \log \frac{1}{\varepsilon} + \log \frac{1}{\delta}\right]$

Recall: $m > \frac{1}{\varepsilon} \ln \frac{|H|}{\delta}$ for finite size H

Intuition for PAC result with infinite H

- d=$VC(H)$
- $m \geq O\left(\frac{1}{\varepsilon}\left[d \log \frac{1}{\varepsilon} + \log \frac{1}{\delta}\right]\right) \sim \frac{1}{\varepsilon}\left[d \log \frac{1}{\varepsilon} + \log \frac{1}{\delta}\right]$
- Finite H: $m > \frac{1}{\varepsilon} \ln \frac{|H|}{\delta}$

$d \log \frac{k}{\varepsilon} \rightarrow \log \frac{k^d}{\varepsilon}$

Can pick h to shatter at most d points in one of two classes

2^d meaningfully different classifiers h: $|H| \sim 2^d$